

# ARTUS – A Framework for Event-based Data Analysis in High Energy Physics

Joram Berger, Fabio Colombo, Raphael Friese, Dominik Haitz, Thomas Hauth, Thomas Müller, Günter Quast, Georg Sieber

Karlsruhe Institute of Technology, Wolfgang-Gaede-Straße 1, D-76131 Karlsruhe

---

## Abstract

ARTUS is an event-based data-processing framework for high energy physics experiments. It is designed for large-scale data analysis in a collaborative environment. The architecture design choices take into account typical challenges and are based on experiences with similar applications. The structure of the framework and its advantages are described. An example use case and performance measurements are presented. The framework is well-tested and successfully used by several analysis groups.

---

## 1. Typical Workflow of High Energy Physics Analyses

The workflow of the experiments collecting data at the Large Hadron Collider (LHC) at CERN is taken as an example to illustrate the organization and needs of a modern High Energy Physics (HEP) analysis. Nonetheless, equal or very similar requirements are also present for HEP experiments not located at the LHC.

Between the data taking at the detector or the generation of simulated events and the presentation of the physical results, the following analysis steps have shown to be the best compromise in terms of turnaround time and efficiency as explained in Figure 1.

The reconstruction of the raw data coming from the experiments is performed on large computing centers organized in the Worldwide LHC Computing Grid (WLCG) [1]. The resulting datasets contain event information and high-level physics objects and are in the order of 100 TB in size.

The data which are relevant for a specific analysis are selected in a step usually referred to as *skimming*. The output of the skimming process is stored in an n-tuple format on local file servers, often owned and operated by a single institute. The size of these files ranges from several 10 GB to few TB.

Finally, the skimmed n-tuples are processed by user code running on local resources. This processing step includes further event selections, event categorization or high-level calculations based on the reconstructed physics objects. The events are usually processed with different configurations, depending on the specific analysis needs. The output is then used for post-processing steps, including the graphical representation of the results and their statistical interpretation.

In many cases, the code for the last step is developed without a predefined strategy or structure, following the evolving needs of the analysis which are often physics-driven and not always entirely predictable in the code design phase. This leads to convoluted, tangled and unstructured code, which is difficult to understand and to maintain and also difficult to hand over to new people joining the analysis effort. Also, code is rarely shared among different groups, although the basic parts of an analysis, such as handling of input data in the same formats, their processing and the selection and identification of physics objects follow similar or equivalent specifications. Looking at the user side, an analysis group usually consists of several people having different academic levels and different programming skills, some of them joining the analysis effort only for a short period of time and limited to a well-defined task.




Figure 1: Characteristic of an event-based analysis in high energy physics, numbers taken from [2]. Data size is reduced by six orders of magnitude, the turnaround time by a factor of 10000. I/O intensive operations are done highly parallelized in the WLCG. The data reduction is achieved by applying analysis specific event selection steps as early as possible. Performing the data-intensive operations only in the first step of this sequence reduces the run time of the analysis.

## 2. Requirements and Challenges for Analysis Software

Based on the workflow described in the previous section, the following demands for an analysis framework arise:

**Performance** The main performance indicator is the *turnaround time* which is the time needed to run the entire analysis in case new requirements arise. Since the runtime is often limited by the transfer speed of the storage servers, the software should load each event only once and perform all necessary processing steps. Still, since analysis requirements sometimes change within a few hours, a high level of flexibility is needed. This explains the need for a modular structure, where only those parts of the code are executed that are really necessary for the analysis.

**Maintainability** Source code maintenance, i.e. to be up-to-date with the latest requirements while not losing backwards compatibility, is a challenge for each software project. To make the entry level as low as possible, beginners can start with an example analysis. Up-to-date recipes should allow the installation and execution of the example analysis out-of-the box. Steering the analysis in detail with configuration files makes it possible to re-use the same code for differing analysis requirements. The modules should encapsulate the actual operations providing and enforcing a clear code structure while the data is stored in the desired data format.

**Reliability** In most cases, analysis code is written from scratch for very specific tasks and it is used only by very small groups of analysts. This leads to frequent re-implementations and low trust in reliability. An ideal framework should allow to re-use analysis code due to its modular structure as well as sharing it between different analysis groups, and therefore improving code quality significantly.

Source code and configuration files are the most precise form of describing the processes. Therefore, readability also for people who have not written the code is crucial. Also, if possible, existing and well-established libraries of the HEP community are interfaced.

**Flexibility** While not aiming at feature-completeness, the software should make it easy to cover all analysis needs by being flexible and extendible. This comprises new data formats, physical quantities and analysis steps, for instance.

## 3. The ARTUS Framework

The framework introduced in this section handles the last step of the workflow, the user analysis.

Following the structure of the CMS software framework [4, 3], the ARTUS framework<sup>1</sup> provides a front-to-end solution for the analysis of event-based data. The basic concept of modules processing events within an event loop is complemented by tools for the configuration of every step. The complete structure is explained in the next section.

This general framework is applicable to any kind of event processing analysis. The well-tested core is shared among various analyses, providing a reliable fundament for any new analysis. The highly modular structure guides the user to develop analysis software that is easily maintainable and extensible. Thus, the ARTUS framework avoids the aforementioned problems of separate user code for individual analyses.

The ARTUS framework is written in C++. The core of the framework only depends on the boost library [5]. Further dependencies may be introduced by the actual analysis code. C++ is chosen for performance reasons and in order to be able to integrate existing frameworks and libraries that are commonly used in high energy physics.

Multi-core techniques are not required since the processing can be trivially parallelized by splitting the input dataset and sending it to different computing instances. This is possible because the processed events are known to be independent of each other. The framework comes with a built-in possibility to submit analysis jobs with the job submission tool grid-control.

#### 4. Structure of the Framework

The basic concept of the framework is taken from the code used for the Z + jet analysis documented in [2]. It is generalized and extended to satisfy the needs of various analysis groups. The structure is illustrated in Figure 2.

ARTUS respects the separation of code and configuration. Settings specify what has to be done and with which parameters. The code performs the selected steps—ideally without requiring the user to be fully aware of implementation details. Hence, the JSON file completely defines the analysis. It contains information on the input data and settings for the pipelines and processors described below. Each of the modules has access to the settings read in from the configuration file. The ARTUS package comes with sophisticated tools written in Python which help in generating this JSON file.

The input is handled by an event provider. Its purpose is to manage the import of event contents from any kind of input to the internal event object and to provide information about the total number of events to be processed.

The events are then sequentially processed by processors organized in pipelines. First, a so-called global pipeline is run. Then the execution is split up into multiple so-called local pipelines which are also processed sequentially. The composition of a pipeline is illustrated in Figure 3.

Each pipeline has its own settings. The local pipelines can access both their own settings and the ones of the global pipeline. Thus, a mechanism exists that enables processing the same event with different configurations, whereas the costly reading of the input is only done once. Similarly, a product object is managed by each pipeline. The product is meant to store new quantities produced in the analysis. The product for the global pipeline is empty at the beginning of the processing of a new event. Local pipelines start with a copy of the global product.

A pipeline consists of an arbitrary number of processor modules that are run in a sequence. Three kinds of processors are distinguished:

**Producers** are meant to calculate new quantities based on the information in the event and the product according to the settings of the pipeline.

**Filters** apply a selection based on event properties. In case an event does not fulfill a requirement defined by a filter, the execution of the subsequent processors can be skipped.

---

<sup>1</sup><https://github.com/artus-analysis/artus>

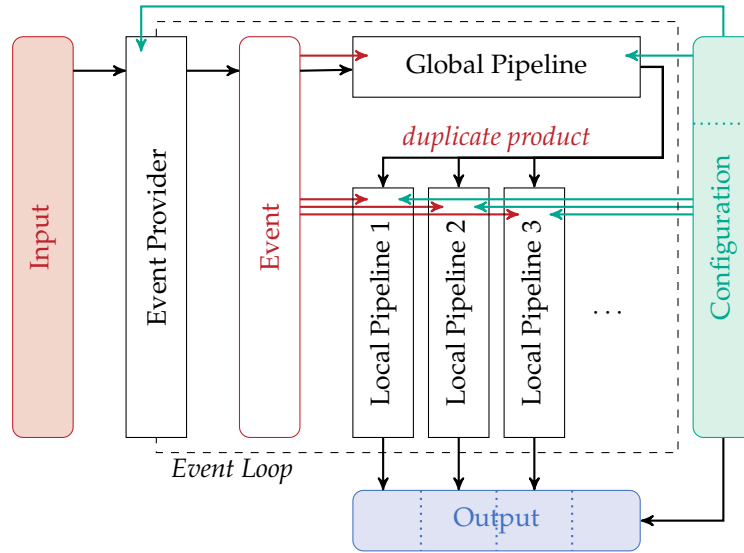


Figure 2: Structure of an ARTUS analysis. The input is read by an event provider. Within the pipelines the event content is analyzed by the processors. Consumers in local pipelines write results to a common output. All parts of the analysis are configurable.

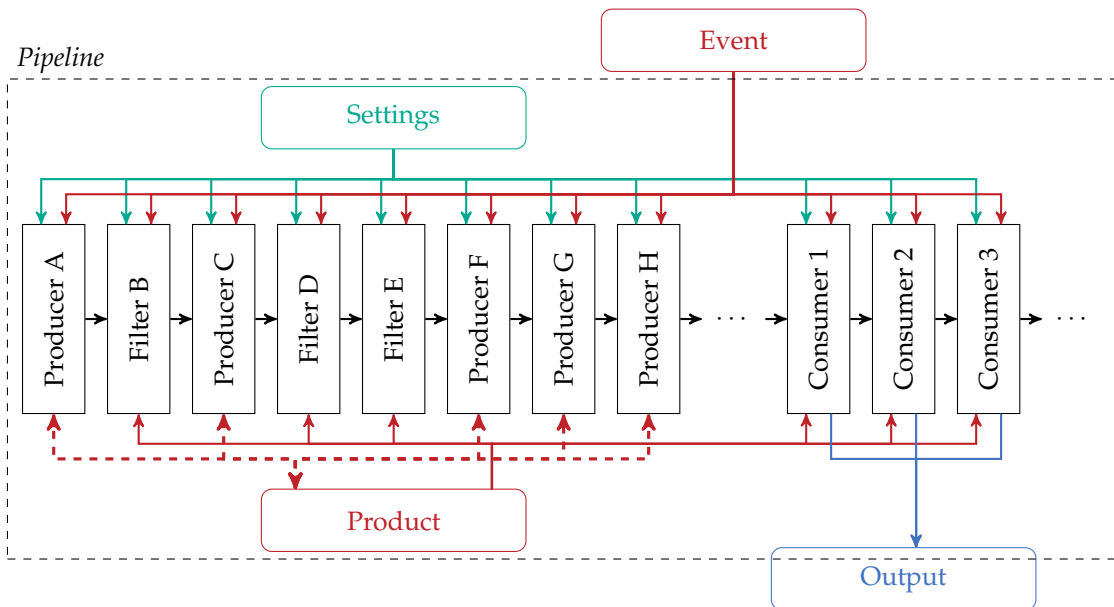


Figure 3: Organization of processors in a pipeline. Producers and filters are run before the consumers. All processors can access the pipeline settings, the event content and the product. Only producers have write access to the product. Consumers can write results to the output file.

**Consumers** can be added to local pipelines. They are executed after the producers and filters. The task of consumers is to take quantities from the event and the product and store them in the output according to the pipeline settings.

Producers and filters can be configured in an arbitrary order. Consumers are executed after the chain of producers and filters. The output together with the complete configuration is written to a ROOT [6] file. These processors are created and managed by a dedicated factory class according to the configuration.

This structure extends the established concepts of existing frameworks [3] to the following four main advantages:

**Modular structure** Analysis steps are implemented in processors, breaking down the complete analysis in well-defined and easily understandable and writable parts.

**Re-utilization of code** Processors can be shared among different analysis groups. Every user contributes to the testing of existing code and improves the confidence in it.

**Configurability** Each module uses settings which are read from the configuration. A separate sub-set of the configuration is created for each pipeline.

**High performance** The concept of pipelines allows to analyze the data with different configurations at the same time, while the data is read in only once and common processing steps are shared in the global pipeline. Therefore, the cost-intensive input operations are reduced to a minimum.

## 5. Building an End-user Analysis

In order to create a custom analysis, the user can build his own code on the ARTUS classes by deriving from them. He must provide an implementation of the settings, the event and the product class. The settings class defines the tags that are read from the configuration file and the types of the corresponding values. The *event* class defines the event content which is read from the input and the *product* class contains members for new quantities calculated during the analysis stage. These classes need to be derived from the base classes provided by the core of the ARTUS framework.

An implementation of an *event provider* must be available. The event provider must know how to read event content from the input files and how to store them in the event object. Furthermore, a factory class defining the creation of analysis-specific processors is required. These classes also have to be derived from the ARTUS base classes.

The most important part of an analysis are the *processors*. The abstract processors of the ARTUS core need to be implemented to realize the physics needs of the analysis. After the basic structure of an analysis is set up, most changes and additions are only related to the processors.

Analyses sharing the input file format can easily also share processors which implement basic analysis steps in a general fashion. Analysis-specific parts can either derive from existing processor classes or be modeled in new processors.

## 6. HARRYPLOTTER– A Python Framework for Post-processing

The computationally expensive event processing steps are followed by the post-processing of the results and their presentation. From a technical point of view, the focus here is rather on flexibility than on performance. The inputs for this step are either histograms filled in the previous step or n-tuples whose entries can be filled into histograms with reasonable computing costs. For this reason, Python has been chosen as language for the HARRYPLOTTER framework<sup>2</sup>. As the name implies, the main focus is on plotting of results. The modular structure allows to integrate multiple analysis steps and several output formats.

The structure of the framework is illustrated in Figure 4. It follows the idea of processors in the ARTUS framework. Three kinds of processor modules are executed subsequently:

---

<sup>2</sup><https://github.com/artus-analysis/Artus/tree/master/HarryPlotter>

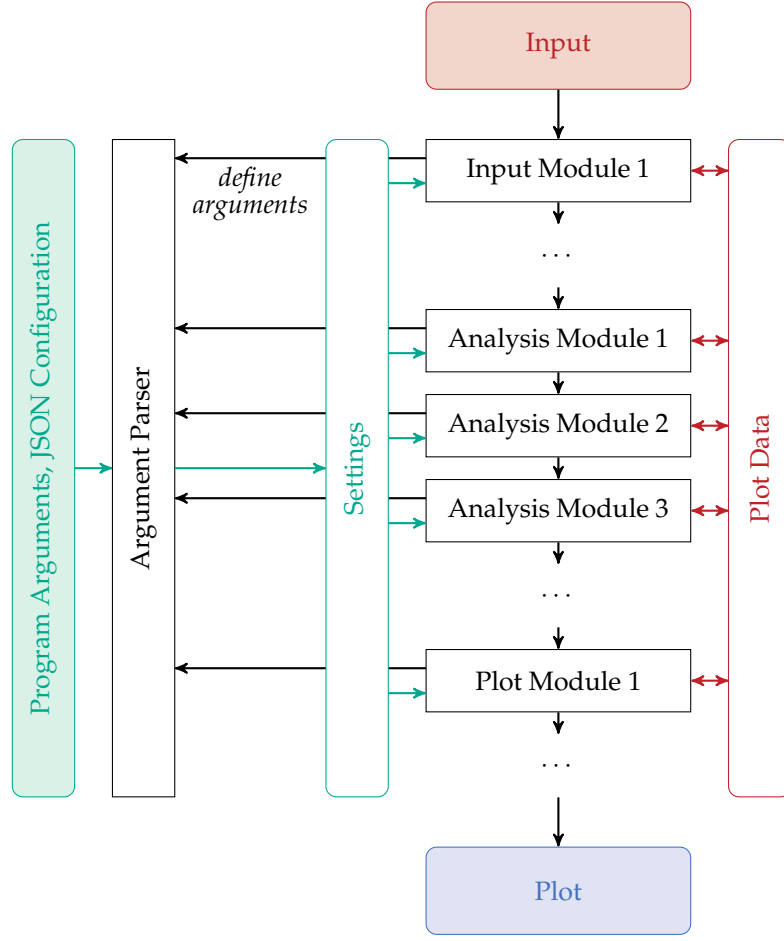


Figure 4: Structure of the HARRYPLOTTER framework. The processing of input modules is followed by the run of analysis modules and concluded by the execution of plot modules. The configuration via a command line interface or via JSON files is parsed by an argument parser whose arguments are defined by the individual modules.

1. One or, in special cases, more input modules are run to read data from an input source. The input format is defined by the input module. The most commonly used input module is the one that reads histograms or other objects from ROOT files.
2. An arbitrary number of analysis modules is appended. They can perform calculations based on the inputs or the results of previously run analysis modules.
3. One or, in special cases, multiple plot modules are executed in the last step. Their task is to output the processed data in the specified format. In most cases, the output is a plot, but it is also possible to store the objects in new ROOT files.

All processors share a common plot data object and possible meta-data where they can add new information and read the existing one. For example, the input modules add histograms read from input files and the plot modules access these histograms.

The configuration is provided by an argument parser that can handle both arguments from a command line interface and content from JSON configuration files or Python dictionaries. This makes it possible to either use the HARRYPLOTTER executable configured via the command line or to call the main function within another Python script. The executable can perform one single plot per run. For the execution within a script, functions exist to create multiple plots in parallel.

The parsing of the arguments is performed in two steps: First, the sequence of modules to be executed

is determined. These modules are then initialized and can define additional arguments needed for their configuration, such that every module can manage its own settings but can also access the settings of the entire program.

## 7. A Real World Application – the $H \rightarrow \tau\tau$ Analysis

Currently, the ARTUS framework is used for various user analyses of CMS data, ranging from jet energy calibration and QCD studies to Higgs boson measurements and searches for physics beyond the Standard Model. As an example, the technical aspects of the  $H \rightarrow \tau\tau$  analysis are explained below.

The  $H \rightarrow \tau\tau$  analysis makes use of skimmed data sets in the  $\kappa$ APPA format. The  $\kappa$ APPA framework [7] is designed for skimming CMS data into small n-tuples that can be analyzed independently of the CMSSW software. The only dependence is the ROOT framework. The n-tuples store high-level information such as four-momentum vectors, collections of leptons and vertices. Similar to the ARTUS framework,  $\kappa$ APPA aims to satisfy general analysis needs.

A large part of the ARTUS software is specialized on the analysis of  $\kappa$ APPA n-tuples<sup>3</sup>. An event provider is available for reading  $\kappa$ APPA n-tuples. The event class contains members for all important physics objects used in analyses. The branches of the n-tuples, that should be read in from the input files, are configurable. This reduces the amount of information transferred from the input files and therefore accelerates the analysis significantly.

The structure of the pipelines is chosen such that a separate pipeline exists for each decay channel. Each of the pipelines is duplicated and modified in order to implement the different settings for shifts of quantities affected by systematic uncertainties. For example, the  $\tau$  energy scale is shifted up and down with respect to the nominal value. The change of the  $\tau$  four-momentum is then propagated through the entire analysis and the effect on the final result is studied. For this, each pipeline for decay channels involving hadronically decaying  $\tau$  leptons is duplicated three times and the setting for the  $\tau$  energy scale shift is modified accordingly. The global pipeline is used to perform actions that are common for all channels. This is for example the identification and correction of jets as well as the filtering of valid data events.

The producers and filters are implemented to perform small tasks of the analysis. Main producers are the ones for the identification and the selection of valid leptons and jets. Further producers perform the matching with trigger or generator objects and perform the splitting into decay channels or the categorization. Also the mass reconstruction is done in a separate producer. The selection of events according to high level trigger decisions as well as the selection of decay channels are examples for filters. The filters are executed as early as possible in the chain of processors in order to avoid running subsequent processors without real need.

Table 1 shows a runtime measurement for the baseline part of the  $H \rightarrow \tau\tau$  in the four main decay channels ( $\mu\tau$ ,  $e\tau$ ,  $e\mu$  and  $\tau\tau$ ). It illustrates the main advantage of the ARTUS structure: sub-analyses with different sets of configurations can be performed in the same run avoiding the multiple I/O overhead when the configurations would be executed sequentially in different runs.

The outputs are flat n-tuples suitable for performing the background modeling step which is done in HARRYPLOTTER. Histograms are read while the final event selection is applied. As an example, the selection on the transverse mass is applied on this level because events from control regions need to be preserved for the background modeling step. For each background process, an analysis module has been developed performing the estimation of the samples. Most of them perform calculations on multiple histograms, defining the yield and the shape of the final quantity in signal and control regions. They add a single histogram to the plot data object containing the final estimation for a given sample. Similarly, ratios comparing the compatibility of data and the simulation for the sub-plots are determined and added this way.

The post-processing step either outputs plots or writes the histograms to ROOT files and passes them over to the limit calculation tools.

---

<sup>3</sup><https://github.com/artus-analysis/Artus/tree/master/KappaAnalysis>

Pipelines	Runtime / min	Event Rate / s <sup>-1</sup>
0	5:12	1470
1 (e $\mu$ )	6:31	1174
1 (e $\tau$ )	7:26	1028
1 ( $\mu\tau$ )	7:14	1052
1 ( $\tau\tau$ )	6:07	1247
2 ( $\mu\tau$ , e $\tau$ )	8:32	896
3 ( $\mu\tau$ , e $\tau$ , e $\mu$ )	9:54	767
4 ( $\mu\tau$ , e $\tau$ , e $\mu$ , $\tau\tau$ )	10:01	758

Table 1: Measurement of the runtime of the baseline  $H \rightarrow \tau\tau$  analysis. The analysis is performed on 10 files with a total size of 6.1 GB containing 452229 simulated  $Z + \text{jet}$  events. The measurement is performed with different numbers of pipelines, where each pipeline implements the analysis of one sub-channel. The values are compared with the execution without any pipelines and without global processors. It is clearly visible, that the time needed for reading the inputs amounts to a large part of the overall runtime and that the simultaneous execution of multiple pipelines can reduce the runtime significantly compared to the case where different configuration are performed sequentially including separate IO operations.

## 8. Conclusion

The ARTUS analysis framework presented in this paper offers a comprehensive and well-integrated solution for event-based data analyses in the HEP domain. In conjunction with the post-processing toolkit HARRYPLOTTER, all typical data analysis needs are covered.

The modular and structured manner, in which the processing required for an analysis, can be expressed within the ARTUS framework, allowed the involved research groups to share and re-use source code on a scale not possible before. The concept underlying the ARTUS framework also enabled sharing program code between analyses and even across different universities and research centers.

Especially new students joining the analysis groups can contribute very quickly. Having understood general concept and structure of an ARTUS analysis, they can contribute by writing their own producers or filters while the full complexity of the event-processing is taken care of by the ARTUS framework. Using the ARTUS framework, it is much less work to set up a new analysis because a rich library of processors, filters and consumers already exists and can be reused for new analyses.

## References

- [1] <http://wlcg-public.web.cern.ch>
- [2] Joram Berger, *Search for the Higgs Boson Produced via Vector-Boson Fusion in the Decay Channel  $H \rightarrow \tau\tau$* , CERN-THESIS-2014-092.
- [3] <https://github.com/cms-sw/cmssw>
- [4] CMS Collaboration, *The CMS Experiment at the CERN LHC*, JINST **3** (2008) S08004, DOI: 10.1088/1748-0221/3/08/S08004,
- [5] <http://www.boost.org>
- [6] Rene Brun and Fons Rademakers, *ROOT – An Object Oriented Data Analysis Framework*, Proceedings AIHENP’96 Workshop, Lausanne, Sep. 1996, Nucl. Inst. & Meth. in Phys. Res. A **389** (1997) 81–86. See also <http://root.cern.ch>.
- [7] <https://github.com/KappaAnalysis/Kappa>